

Section 1: Nuclear and magnetic model comparison

Figure E1 shows the best fit models for the monolayer (a-c), bilayer (d-f), trilayer (g-i), and graded (j-l) samples, determined from fits to PNR data. The nuclear scattering length density is defined as

$$\rho = \sum_i N_i b_i,$$

where N is the number density, b is the characteristic nuclear scattering length of a given isotope, and the summation is over every type of constituent isotope. The nuclear profiles are indicative of the sample composition, and are shown in the top row of panels in Figure E1. The Pd cap, Co/Pd multilayer, Pd seed, native oxide layer, and Si substrate are delineated by variations in ρ . Data for the monolayer sample are fit well by a model with uniform average ρ throughout the Co/Pd multilayer (a). For the bilayer (d), trilayer (g), and graded (j) samples, increased t_{Co} in the Co/Pd multilayer is manifest as a reduced ρ near the Pd cap ($b_{Co} = 2.49$ fm, $b_{Pd} = 5.91$ fm). Field dependent magnetization depth profiles are shown in the center row of panels. The monolayer sample data is fit well by a uniform magnetization model (b), while non-uniform magnetization profiles are required to fit the bilayer (e), trilayer (h), and graded (k) sample data. However, from these profiles, it is unclear whether or not the observed magnetization gradients are partly due to graded anisotropy, or if the gradients are wholly due to the increased Co content for the thicker Co layers near the surface. To distinguish between these two causes, the bottom row of panels in Figure E1 shows the magnetization profiles normalized by the respective maximum values of each profile. If the normalized profiles show no field

dependence, it is indicative of zero anisotropy gradient. Conversely, a field-dependent normalized magnetization profile indicates that magnetic moments at different depths in of the Co/Pd stack respond to field at different rates – i. e. the magnetic anisotropy varies throughout the stack. The bilayer (f), trilayer (i), and graded (l) samples all exhibit such field dependent normalized magnetic depth profiles – thus, these samples all exhibit depth dependent anisotropy.

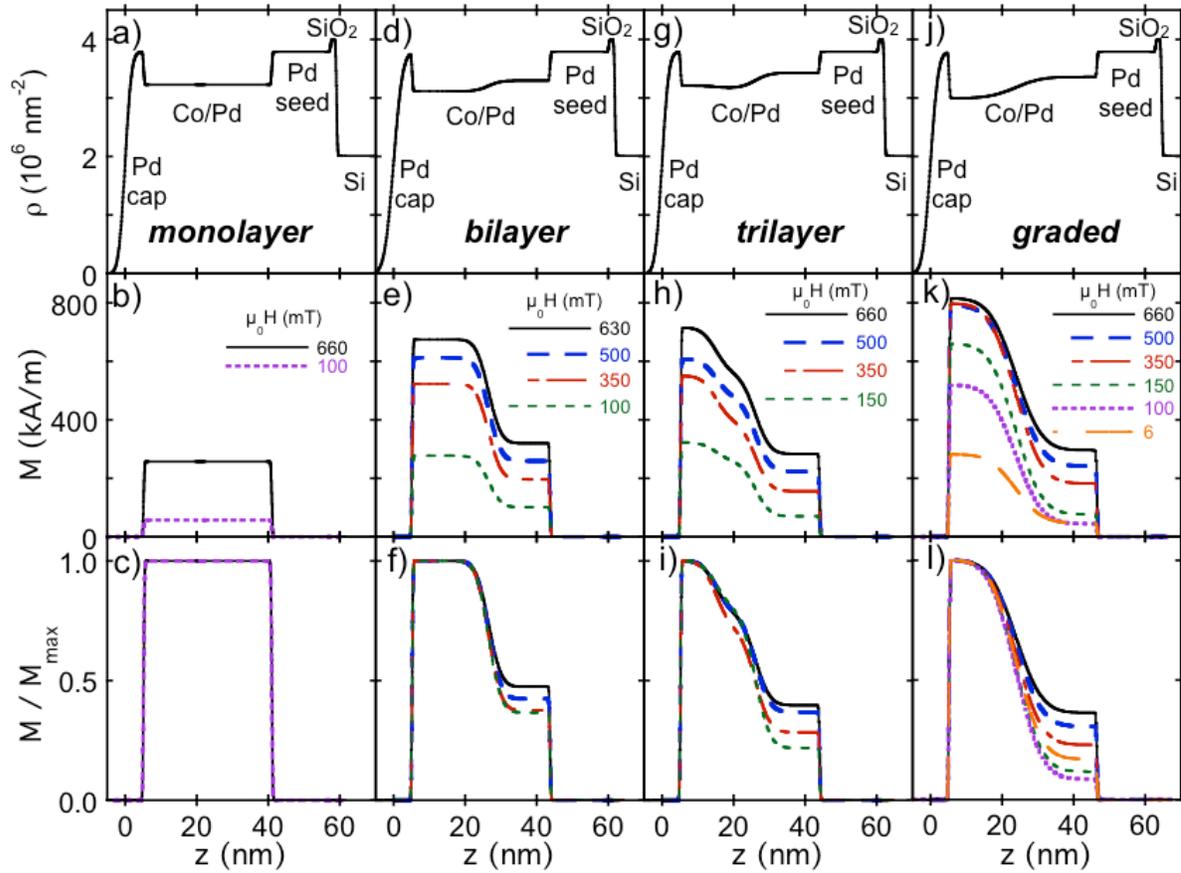


Figure E1: Field-dependent model profiles for the monolayer (a-c), bilayer (d-f), trilayer (g-i), and graded (j-l) samples. The normalized magnetization profiles for the bilayer (f), trilayer (i), and graded (l) samples are field dependent – thus, the samples exhibit depth-dependent anisotropy.

Section 2: Comparison of graded and uniform anisotropy models.

As described in the regular manuscript, the model profile corresponding to the best fit to the graded sample PNR data reveals a graded anisotropy profile. To confirm our sensitivity to this result, we compared the best-fit model (Figure E2, left column) to a constant anisotropy model (Figure E2, right column). For the constant anisotropy model, the fitting procedure and fitting parameters were the same as for the best-fit graded anisotropy model, except that the magnetization profiles corresponding to different values of H were required to be proportional to one another. Put another way, the magnetic profiles normalized by the respective maximum magnetization values were constrained to be identical. Thus, the constrained model allowed for variations in depth-dependent magnetization due to variations in total Co concentration (i.e. variations in Co layer thickness throughout the Co/Pd stack), but did not allow for variations in magnetization due to variations in anisotropy. PNR data for the graded sample fitted using the graded anisotropy model, is compared to the same data fitted using the constant anisotropy model in Figures E3-E8. PNR data and fits are shown multiplied by Q^4 for clarity. To highlight differences between the fits and the data, fitting residuals (the difference between the data and the fit divided by the size of the error bar) are shown directly beneath the plots of the fitted data. Residuals are a useful way to gauge fit quality, as a residual absolute value larger than 1 indicates a deviation of the fit from the data larger than can be expected from the uncertainty in the data. Figures E3-E8 clearly shows that the graded anisotropy model fits the data extremely well, and that the constant anisotropy model fits the data significantly worse, evidenced by an increased goodness of

fit parameter chi-squared for all but the 6 mT data set (Fig. E-8) for which the two models fit the data essentially equally well. The difference in fitting quality between the two models is particularly noticeable for the 660 mT (Fig. E3) and 150 mT (Fig. E6) data, where residuals for the constant anisotropy model show deviations up to an order of magnitude larger than what would be expected from data uncertainty. This comparison conclusively demonstrates our sensitivity to the sample's anisotropy profile, and confirms the graded anisotropy determined by the best fit.

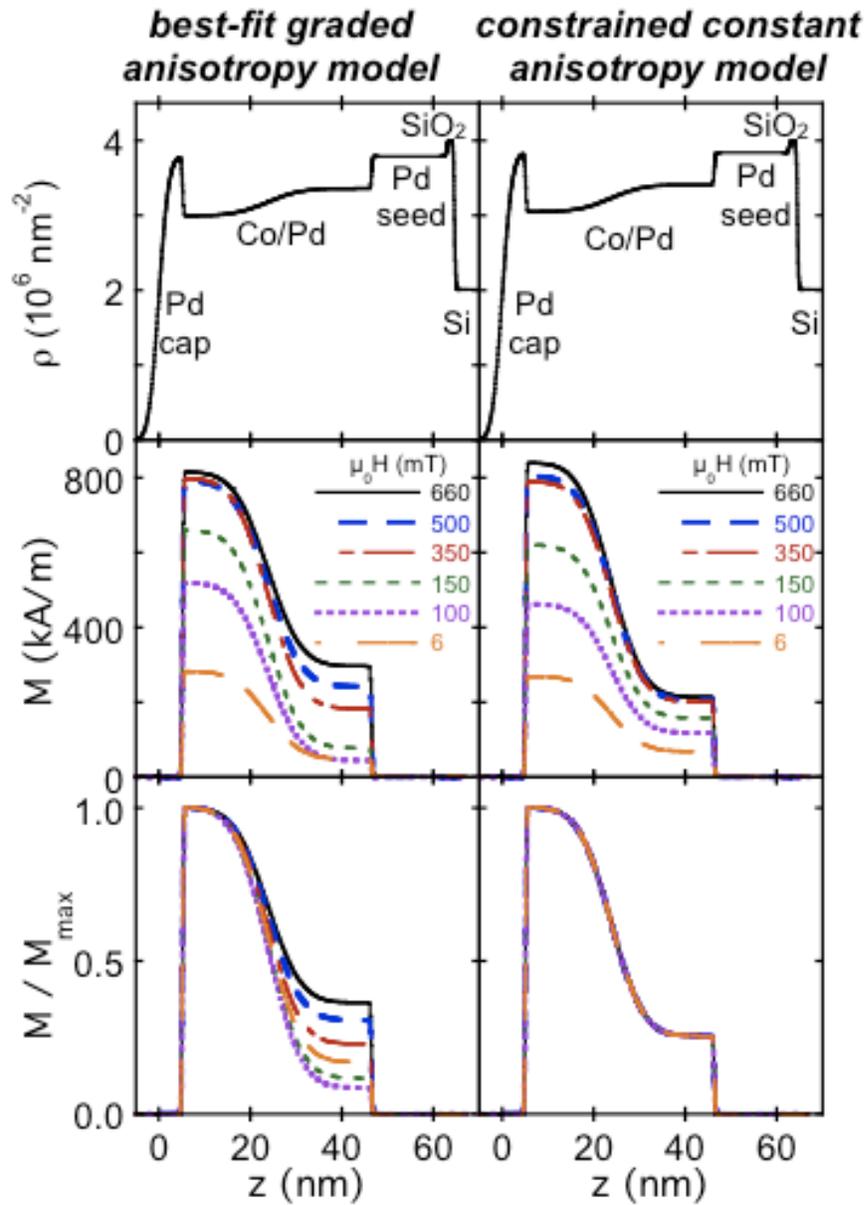


Figure E2: Competing models used to fit the graded sample PNR data. The best-fit graded anisotropy model is shown to the left, and the constrained constant anisotropy model is shown at right.

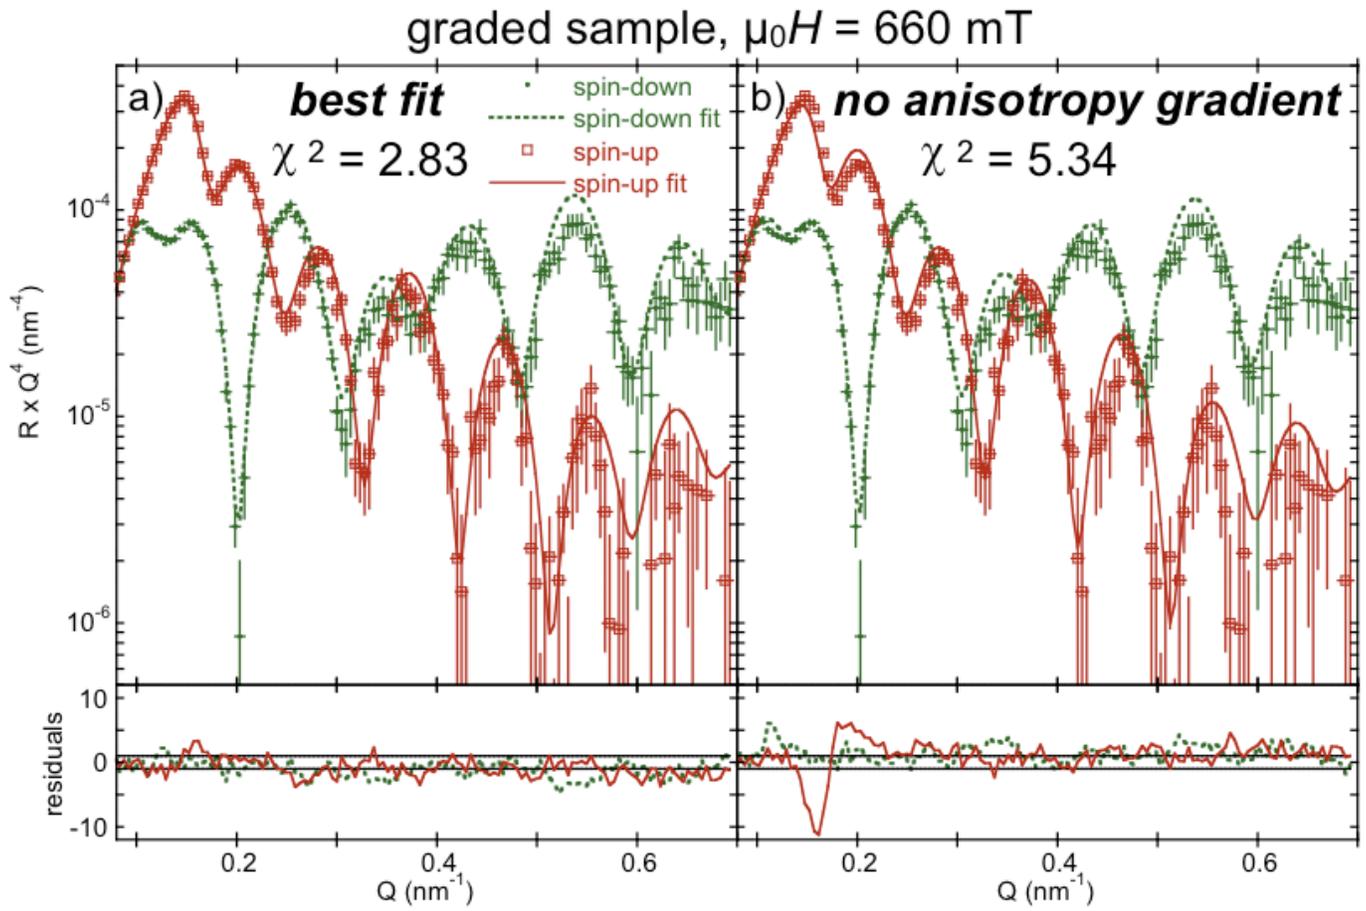


Figure E3: Comparison of graded and constant anisotropy model fits for the 650 mT graded sample data.

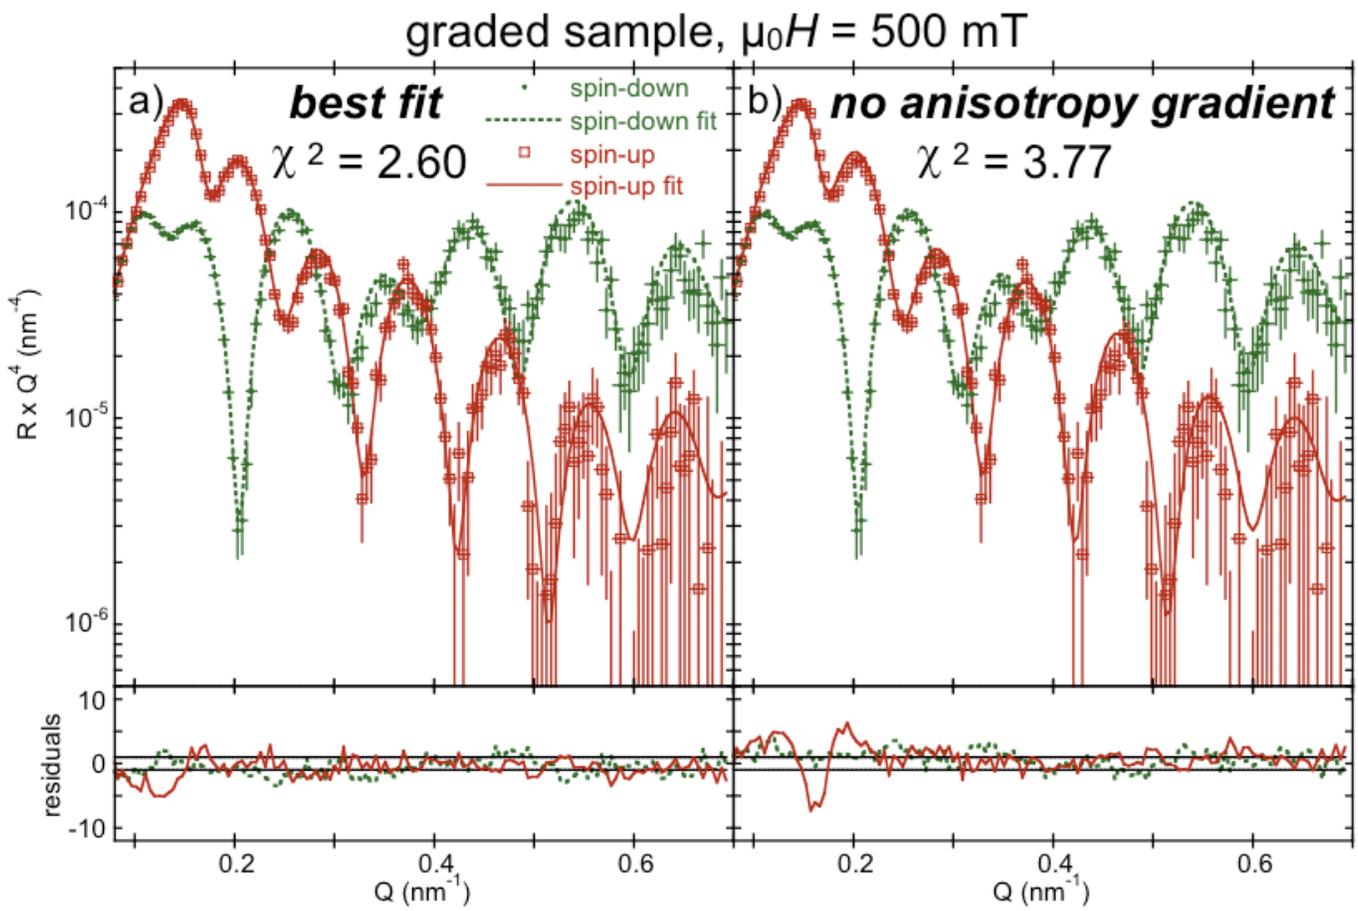


Figure E4: Comparison of graded and constant anisotropy model fits for the 500 mT graded sample data.

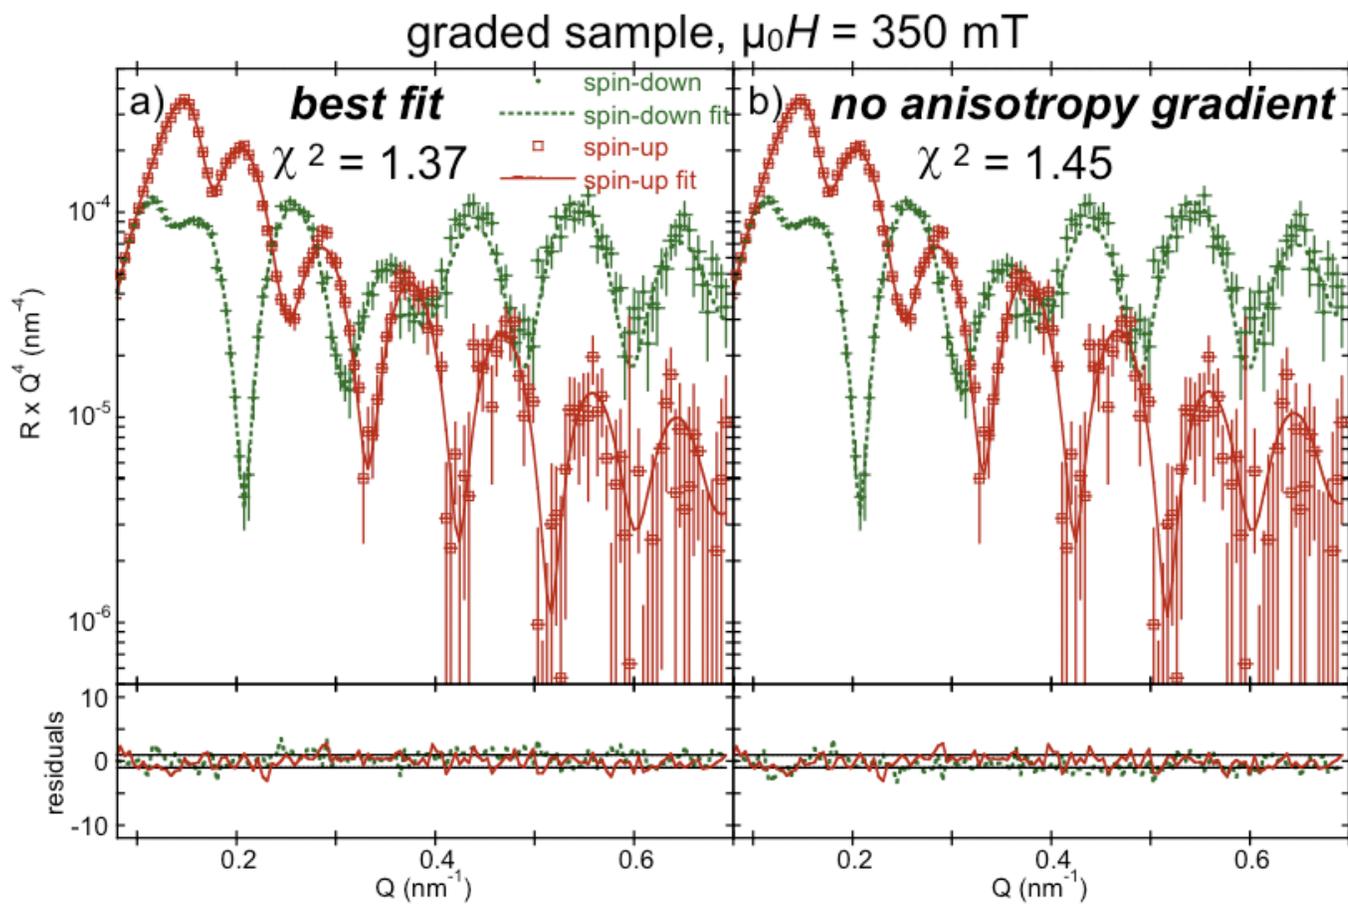


Figure E5: Comparison of graded and constant anisotropy model fits for the 350 mT graded sample data.

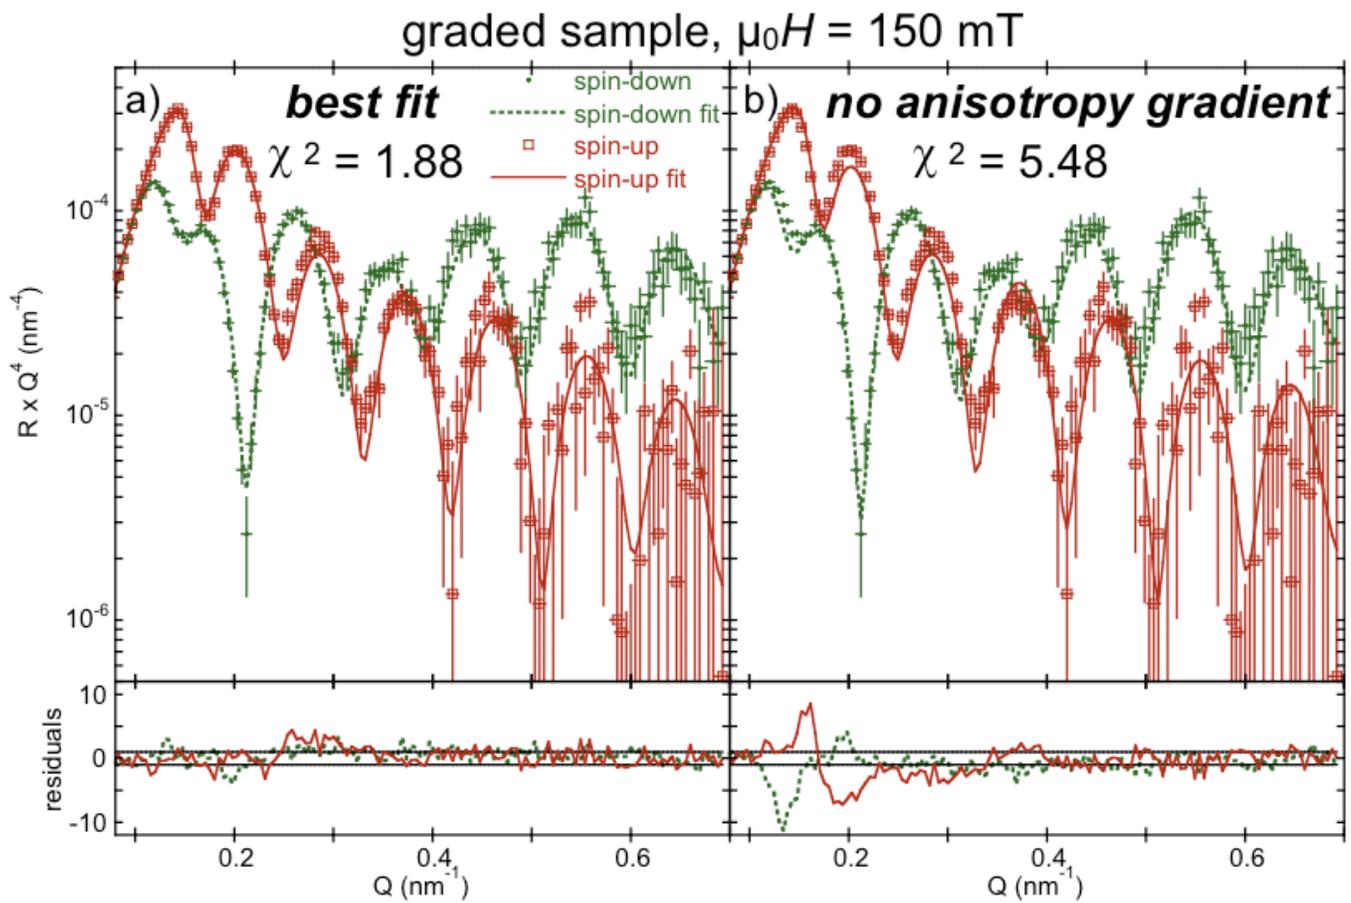


Figure E6: Comparison of graded and constant anisotropy model fits for the 150 mT graded sample data.

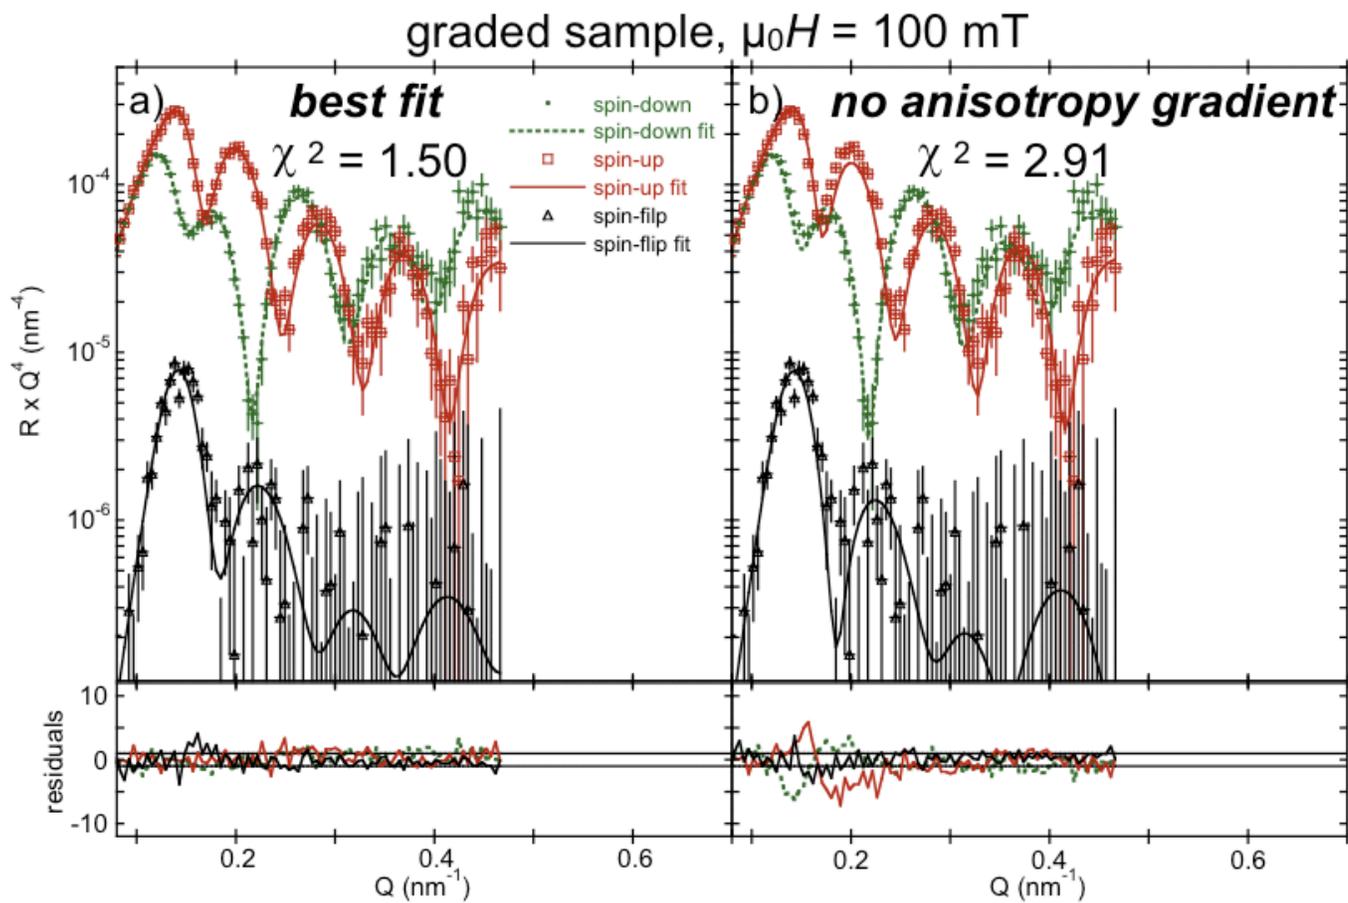


Figure E7: Comparison of graded and constant anisotropy model fits for the 100 mT graded sample data.

graded sample, $\mu_0 H = 6$ mT

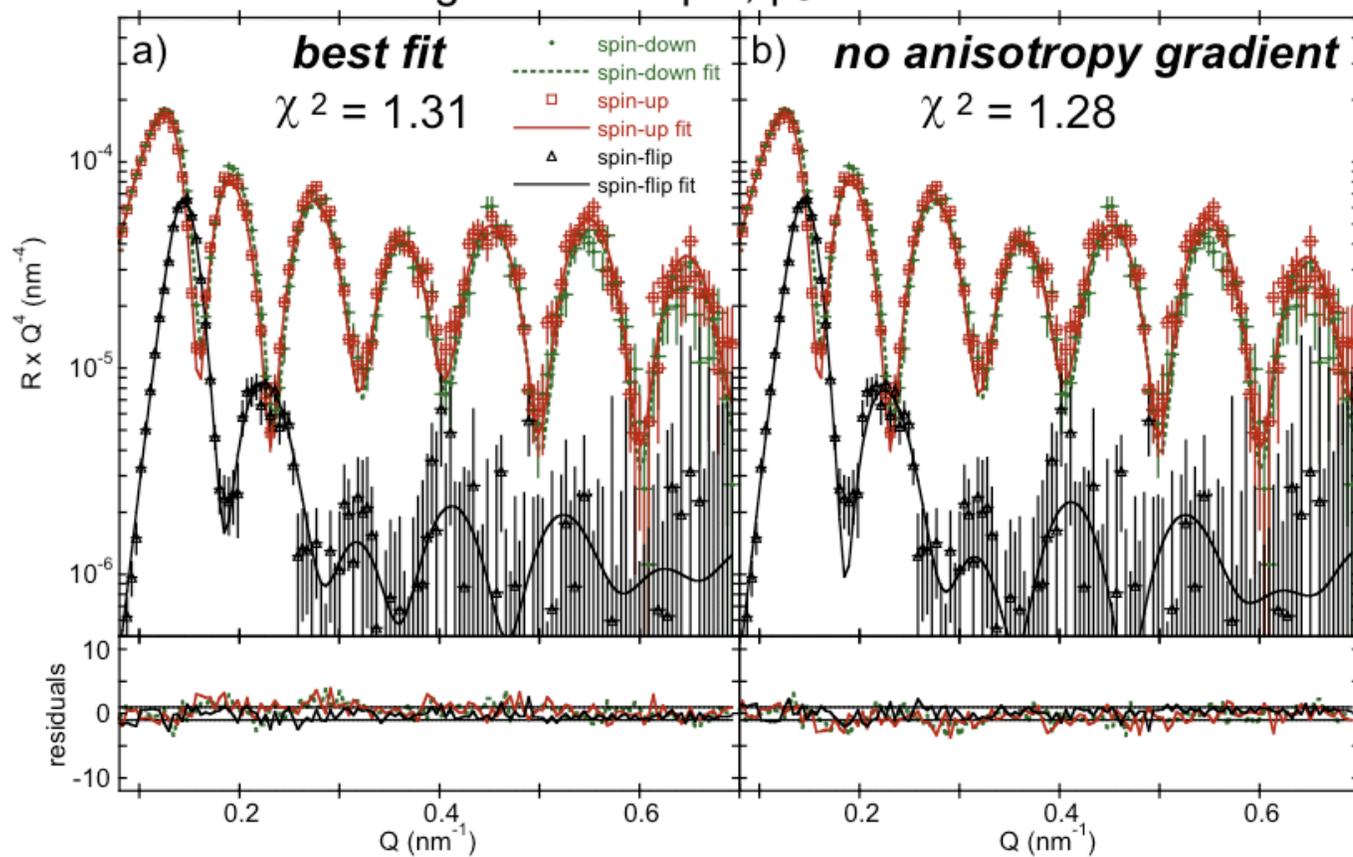


Figure E8: Comparison of graded and constant anisotropy model fits for the 6 mT graded sample data.

Section 3: First order reversal curve (FORC) distributions.

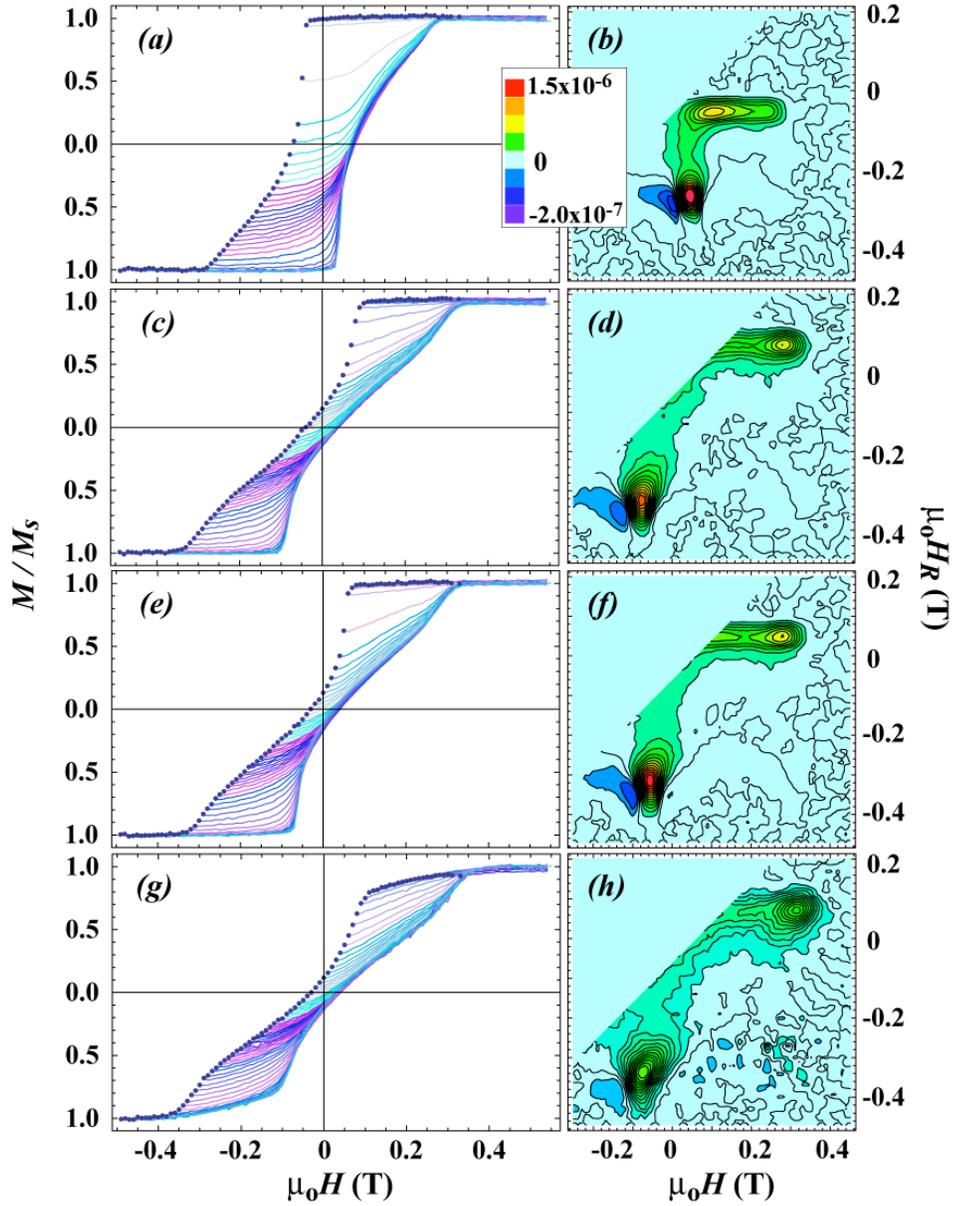


Figure E9– FORCs (left) and corresponding FORC distributions (right) for the hard-layer [(a) – (b)], bilayer [(c) – (d)], trilayer [(e) – (f)], and graded [(g) – (h)] samples. The contour legend is inset into (a) with all contour plots scaled to the same color-grading. The general pattern emerging in the FORC distributions is characteristic of reversal by a domain nucleation, expansion and annihilation process.